arXiv:1402.6763v1 [math.OC] 27 Feb 2014

Linear Programming for Large-Scale Markov Decision Problems

Yasin Abbasi-Yadkori
Queensland University of Technology
yasin.abbasiyadkori@qut.edu.au

Peter L. Bartlett
UC Berkeley and QUT
bartlett@eecs.berkeley.edu

Alan Malek
UC Berkeley
malek@eecs.berkeley.edu

February 28, 2014

Abstract

We consider the problem of controlling a Markov decision process (MDP) with a large state space, so as to minimize average cost. Since it is intractable to compete with the optimal policy for large scale problems, we pursue the more modest goal of competing with a low-dimensional family of policies. We use the dual linear programming formulation of the MDP average cost problem, in which the variable is a stationary distribution over state-action pairs, and we consider a neighborhood of a low-dimensional subset of the set of stationary distributions (defined in terms of state-action features) as the comparison class. We propose two techniques, one based on stochastic convex optimization, and one based on constraint sampling. In both cases, we give bounds that show that the performance of our algorithms approaches the best achievable by any policy in the comparison class. Most importantly, these results depend on the size of the comparison class, but not on the size of the state space. Preliminary experiments show the effectiveness of the proposed algorithms in a queuing application.

1 Introduction

We study the average loss Markov decision process problem. The problem is well-understood when the state and action spaces are small (Bertsekas, 2007). Dynamic programming (DP) algorithms, such as value iteration (Bellman, 1957) and policy iteration (Howard, 1960), are standard techniques to compute the optimal policy. In large state space problems, exact DP is not feasible as the computational complexity scales quadratically with the number of states.

A popular approach to large-scale problems is to restrict the search to the linear span of a small number of features. The objective is to compete with the best solution within this comparison class. Two popular methods are Approximate Dynamic Programming (ADP) and Approximate Linear

Programming (ALP). This paper focuses on ALP. For a survey on theoretical results for ADP see (Bertsekas and Tsitsiklis, 1996, Sutton and Barto, 1998), (Bertsekas, 2007, Vol. 2, Chapter 6), and more recent papers (Sutton et al., 2009b,a, Maei et al., 2009, 2010).

Our aim is to develop methods that find policies with performance guaranteed to be close to the best in the comparison class but with computational complexity that does not grow with the size of the state space. All prior work on ALP either scales badly or requires access to samples from a distribution that depends on the optimal policy.

This paper proposes a new algorithm to solve the Approximate Linear Programming problem that is computationally efficient and does not require knowledge of the optimal policy. In particular, we introduce new proof techniques and tools for average cost MDP problems and use these techniques to derive a reduction to stochastic convex optimization with accompanying error bounds. We also propose a constraint sampling technique and obtain performance guarantees under an additional assumption on the choice of features.

1.1 Notation

Let X and A be positive integers. Let $\mathcal{X} = \{1, 2, \ldots, X\}$ and $\mathcal{A} = \{1, 2, \ldots, A\}$ be state and action spaces, respectively. Let Δ_S denote probability distributions over set S. A policy π is a map from the state space to $\Delta_{\mathcal{A}}$: $\pi : \mathcal{X} \to \Delta_{\mathcal{A}}$. We use $\pi(a|x)$ to denote the probability of choosing action a in state x under policy π. A transition probability kernel (or transition kernel) $P : \mathcal{X} \times \mathcal{A} \to \Delta_{\mathcal{X}}$ maps from the direct product of the state and action spaces to $\Delta_{\mathcal{X}}$. Let P^π denote the probability transition kernel under policy π. A loss function is a bounded real-valued function over state and action spaces, $\ell : \mathcal{X} \times \mathcal{A} \to [0, 1]$. Let $M_{i,:}$ and $M_{:,j}$ denote ith row and jth column of matrix M respectively. Let $\|v\|_{1,c} = \sum_i c_i |v_i|$ and $\|v\|_{\infty,c} = \max_i c_i |v_i|$ for a positive vector c. We use $\mathbf{1}$ and $\mathbf{0}$ to denote vectors with all elements equal to one and zero, respectively. We use $\wedge$ and $\vee$ to denote the minimum and the maximum, respectively. For vectors v and w, $v \leq w$ means element-wise inequality, i.e. $v_i \leq w_i$ for all i.

1.2 Linear Programming Approach to Markov Decision Problems

Under certain assumptions, there exist a scalar λ_* and a vector $h_* \in \mathbb{R}^X$ that satisfy the Bellman optimality equations for average loss problems,

$$\lambda_* + h_*(x) = \min_{a \in \mathcal{A}} \left[\ell(x, a) + \sum_{x' \in \mathcal{X}} P_{(x,a),x'} h_*(x') \right] .$$

The scalar λ_* is called the optimal average loss, while the vector h_* is called a differential value function. The action that minimizes the right-hand side of the above equation is the optimal action in state x and is denoted by $a_*(x)$. The optimal policy is defined by $\pi_*(a_*(x)|x) = 1$. Given ℓ and

P, the objective of the *planner* is to compute the optimal action in all states, or equivalently, to find the optimal policy.

The MDP problem can also be stated in the LP formulation (Manne, 1960),

$$\begin{aligned} &\max_{\lambda,h} \lambda\,, \\ &\text{s.t.} \quad B(\lambda \mathbf{1} + h) \le \ell + Ph\,, \end{aligned} \tag{1}$$

where $B \in \{0,1\}^{XA \times X}$ is a binary matrix such that the ith column has A ones in rows $1+(i-1)A$ to iA. Let v_π be the stationary distribution under policy π and let $\mu_\pi(x,a) = v_\pi(x)\pi(a|x)$. We can write

$$\begin{aligned} \pi_* &= \operatorname*{argmin}_\pi \sum_{x\in\mathcal{X}} v_\pi(x) \sum_{a\in\mathcal{A}} \pi(a|x)\ell(x,a) \\ &= \operatorname*{argmin}_\pi \sum_{(x,a)\in\mathcal{X}\times\mathcal{A}} \mu_\pi(x,a)\ell(x,a) = \operatorname*{argmin}_\pi \mu_\pi^\top \ell\,. \end{aligned}$$

In fact, the dual of LP (1) has the form of

$$\begin{aligned} &\min_{\mu\in\mathbb{R}^{XA}} \mu^\top \ell\,, \\ &\text{s.t.} \quad \mu^\top \mathbf{1} = 1,\, \mu \ge \mathbf{0},\, \mu^\top (P - B) = \mathbf{0}\,. \end{aligned} \tag{2}$$

The objective function, $\mu^\top \ell$, is the average loss under stationary distribution μ. The first two constraints ensure that μ is a probability distribution over state-action space, while the last constraint ensures that μ is a stationary distribution. Given a solution μ, we can obtain a policy via $\pi(a|x) = \mu(x,a)/\sum_{a'\in\mathcal{A}} \mu(x,a')$.

1.3 Approximations for Large State Spaces

The LP formulations (1) and (2) are not practical for large scale problems as the number of variables and constraints grows linearly with the number of states. Schweitzer and Seidmann (1985) propose approximate linear programming (ALP) formulations. These methods were later improved by de Farias and Van Roy (2003a,b), Hauskrecht and Kveton (2003), Guestrin et al. (2004), Petrik and Zilberstein (2009), Desai et al. (2012). As noted by Desai et al. (2012), the prior work on ALP either requires access to samples from a distribution that depends on optimal policy or assumes the ability to solve an LP with as many constraints as states. (See Section 2 for a more detailed discussion.) Our objective is to design algorithms for very large MDPs that do not require knowledge of the optimal policy.

In contrast to the aforementioned methods, which solve the primal ALPs (with value functions as variables), we work with the dual form (2) (with stationary distributions as variables). Analogous

to ALPs, we control the complexity by limiting our search to a linear subspace defined by a small number of *features*. Let d be the number of features and Φ be a $(XA) \times d$ matrix with features as column vectors. By adding the constraint $\mu = \Phi\theta$, we get

$$\begin{aligned} &\min_{\theta} \theta^\top \Phi^\top \ell \,, \\ &\text{s.t.} \quad \theta^\top \Phi^\top \mathbf{1} = 1,\, \Phi\theta \geq \mathbf{0},\, \theta^\top \Phi^\top (P - B) = \mathbf{0} \,. \end{aligned}$$

If a stationary distribution μ_0 is known, it can be added to the linear span to get the ALP

$$\begin{aligned} &\min_{\theta} (\mu_0 + \Phi\theta)^\top \ell \,, \\ &\text{s.t.} \quad (\mu_0 + \Phi\theta)^\top \mathbf{1} = 1,\, \mu_0 + \Phi\theta \geq \mathbf{0},\, (\mu_0 + \Phi\theta)^\top (P - B) = \mathbf{0} \,. \end{aligned} \tag{3}$$

Although $\mu_0 + \Phi\theta$ might not be a stationary distribution, it still defines a policy[1]

$$\pi_\theta(a|x) = \frac{[\mu_0(x,a) + \Phi_{(x,a),:}\theta]_+}{\sum_{a'} [\mu_0(x,a') + \Phi_{(x,a'),:}\theta]_+} \,, \tag{4}$$

We denote the stationary distribution of this policy μ_θ which is only equal to $\mu_0 + \Phi\theta$ if θ is in the feasible set.

1.4 Problem definition

With the above notation, we can now be explicit about the problem we are solving.

Definition 1 (Efficient Large-Scale Dual ALP)**.** *For an MDP specified by ℓ and P, the efficient large-scale dual ALP problem is to produce parameters $\widehat{\theta}$ such*

$$\mu_{\widehat{\theta}}^\top \ell \leq \min \left\{ \mu_\theta^\top \ell : \theta \text{ feasible for } (3) \right\} + O(\epsilon) \tag{5}$$

in time polynomial in d and $1/\epsilon$. The model of computation allows access to arbitrary entries of Φ, ℓ, P, μ_0, $P^\top\Phi$, and $\mathbf{1}^\top\Phi$ in unit time.

Importantly, the computational complexity cannot scale with X and we do not assume any knowledge of the optimal policy. In fact, as we shall see, we solve a harder problem, which we define as follows.

Definition 2 (Expanded Efficient Large-Scale Dual ALP)**.** *Let $V : \Re^d \to \Re_+$ be some "violation function" that represents how far $\mu_0 + \Phi\theta$ is from a valid stationary distribution, satisfying $V(\theta) = 0$ if θ is a feasible point for the ALP (3). The expanded efficient large-scale dual ALP problem is to*

[1] We use the notation $[v]_- = v \wedge 0$ and $[v]_+ = v \vee 0$.

produce parameters $\widehat{\theta}$ *such that*

$$\mu_{\widehat{\theta}}^{\top}\ell \leq \min\left\{\mu_{\theta}^{\top}\ell + \frac{1}{\epsilon}V(\theta) : \theta \in \Re^{d}\right\} + O(\epsilon), \tag{6}$$

in time polynomial in d *and* $1/\epsilon$*, under the same model of computation as in Definition 1.*

Note that the expanded problem is strictly more general as guarantee (6) implies guarantee (5). Also, many feature vectors Φ may not admit any feasible points. In this case, the dual ALP problem is trivial, but the expanded problem is still meaningful.

Having access to arbitrary entries of the quantities in Definition 1 arises naturally in many situations. In many cases, entries of $P^{\top}\Phi$ are easy to compute. For example, suppose that for any state x' there are a small number of state-action pairs (x, a) such that $P(x'|x, a) > 0$. Consider Tetris; although the number of board configurations is large, each state has a small number of possible neighbors. Dynamics specified by graphical models with small connectivity also satisfy this constraint. Computing entries of $P^{\top}\Phi$ is also feasible given reasonable features. If a feature φ_i is a stationary distribution, then $P^{\top}\varphi_i = B^{\top}\varphi_i$. Otherwise, it is our prerogative to design sparse feature vectors, hence making the multiplication easy. We shall see an example of this setting later.

1.5 Our Contributions

In this paper, we introduce an algorithm that solves the expanded efficient large-scale dual ALP problem under a (standard) assumption that any policy converges quickly to its stationary distribution.

Our algorithm take as input a constant S and an error tolerance ϵ, and has access to the various products listed in Definition 1. Define $\Theta = \{\theta \ : \ \theta^{\top}\Phi^{\top}\mathbf{1} = 1 - \mu_0^{\top}\mathbf{1},\ \|\theta\| \leq S\}$. If no stationary distribution is known, we can simply choose $\mu_0 = \mathbf{0}$. The algorithm is based on stochastic convex optimization. We prove that for any $\delta \in (0, 1)$, after $O(1/\epsilon^4)$ steps of gradient descent, the algorithm finds a vector $\widehat{\theta} \in \Theta$ such that, with probability at least $1 - \delta$,

$$\mu_{\widehat{\theta}}^{\top}\ell \leq \mu_{\theta}^{\top}\ell + \frac{1}{\epsilon}\left\|[\mu_0 + \Phi\theta]_{-}\right\|_1 + \frac{1}{\epsilon}\left\|(P - B)^{\top}(\mu_0 + \Phi\theta)\right\|_1 + O(\epsilon\log(1/\delta))$$

holds for all $\theta \in \Theta$; i.e., we solve the expanded problem for $V(\theta)$ equal to the L_1 error of the violation. The second and third terms are zero for feasible points (points in the intersection of the feasible set of LP (2) and the span of the features). For points outside the feasible set, these terms measure the extent of constraint violations for the vector $\mu_0 + \Phi\theta$, which indicate how well stationary distributions can be represented by the chosen features.

The above performance bound scales with $1/\epsilon$ that can be large when the feasible set is empty and ϵ is very small. We propose a second approach and show that this dependence can be eliminated if we have extra information about the MDP. Our second approach is based on the constraint

sampling method that is already applied to large-scale linear programming and MDP problems (de Farias and Van Roy, 2004, Calafiore and Campi, 2005, Campi and Garatti, 2008). We obtain performance bounds, but under the condition that a suitable function that controls the size of constraint violations is available. Our proof technique is different from previous work and gives stronger performance bounds.

Our constraint sampling method takes two extra inputs: functions v_1 and v_2 that specify the amount of constraint violations that we tolerate at each state-action pair. The algorithm samples $O(\frac{d}{\epsilon}\log\frac{1}{\epsilon})$ constraints and solves the sampled LP problem. Let $\widetilde{\theta}$ denote the solution of the sampled ALP and θ_* denote the solution of the full ALP subject to constraints v_1 and v_2. We prove that with high probability,

$$\ell^\top \mu_{\widetilde{\theta}} \leq \ell^\top \mu_{\theta_*} + O(\epsilon + \|v_1\|_1 + \|v_2\|_1) \,.$$

2 Related Work

de Farias and Van Roy (2003a) study the *discounted* version of the primal form (1). Let $c \in \mathbb{R}^X$ be a vector with positive components and $\gamma \in (0,1)$ be a discount factor. Let $L : \mathbb{R}^X \to \mathbb{R}^X$ be the Bellman operator defined by $(LJ)(x) = \min_{a\in\mathcal{A}}(\ell(x,a) + \gamma \sum_{x'\in\mathcal{X}} P_{(x,a),x'} J(x'))$ for $x \in \mathcal{X}$. Let $\Psi \in \mathbb{R}^{X\times d}$ be a feature matrix. The exact and approximate LP problems are as follows:

$$\begin{array}{ll} \max\limits_{J\in\mathbb{R}^X} c^\top J \,, & \max\limits_{w\in\mathbb{R}^d} c^\top \Psi w \,, \\ \text{s.t.} \quad LJ \geq J \,, & \text{s.t.} \quad L\Psi w \geq \Psi w \,. \end{array}$$

which can also be written as

$$\begin{array}{ll} \max\limits_{J\in\mathbb{R}^X} c^\top J \,, & \max\limits_{w\in\mathbb{R}^d} c^\top \Psi w \,, \\ \text{s.t.} \quad \forall (x,a),\, \ell(x,a) + \gamma P_{(x,a),:} J \geq J(x) \,, & \text{s.t.} \quad \forall (x,a),\, \ell(x,a) + \gamma P_{(x,a),:} \Psi w \geq (\Psi w)(x) \,. \end{array} \tag{7}$$

The optimization problem on the RHS is an approximate LP with the choice of $J = \Psi w$. Let $J_\pi(x) = \mathbb{E}\left[\sum_{t=0}^\infty \gamma^t \ell(x_t, \pi(x_t)) | x_0 = x\right]$ be value of policy π, J_* be the solution of LHS, and $\pi_J(x) = \text{argmin}_{a\in\mathcal{A}}(\ell(x,a) + \gamma P_{(x,a),:} J)$ be the greedy policy with respect to J. Let $\nu \in \Delta_\mathcal{X}$ be a probability distribution and define $\mu_{\pi,\nu} = (1-\gamma)\nu^\top (I - \gamma P^\pi)^{-1}$. de Farias and Van Roy (2003a) prove that for any J satisfying the constraints of the LHS of (7),

$$\|J_{\pi_J} - J_*\|_{1,\nu} \leq \frac{1}{1-\gamma} \|J - J_*\|_{1,\mu_{\pi_J,\nu}} \,. \tag{8}$$

Define $\beta_u = \gamma \max_{x,a} \sum_{x'} P_{(x,a),x'} u(x')/u(x)$. Let $U = \{u \in \mathbb{R}^X : u \geq \mathbf{1}, u \in \text{span}(\Psi), \beta_u < 1\}$. Let

w_* be the solution of ALP. de Farias and Van Roy (2003a) show that for any $u \in U$,

$$\|J_* - \Psi w_*\|_{1,c} \leq \frac{2c^\top u}{1-\beta_u} \min_w \|J_* - \Psi w\|_{\infty,1/u} \ . \tag{9}$$

This result has a number of limitations. First, solving ALP can be computationally expensive as the number of constraints is large. Second, it assumes that the feasible set of ALP is non-empty. Finally, Inequality (8) implies that $c = \mu_{\pi_{\Psi w_*},\nu}$ is an appropriate choice to obtain performance bounds. However, w_* itself is function of c and is not known before solving ALP.

de Farias and Van Roy (2004) propose a computationally efficient algorithm that is based on a constraint sampling technique. The idea is to sample a relatively small number of constraints and solve the resulting LP. Let $\mathcal{N} \subset \mathbb{R}^d$ be a known set that contains w_* (solution of ALP). Let $\mu^V_{\pi,c}(x) = \mu_{\pi,c}(x)V(x)/(\mu^\top_{\pi,c}V)$ and define the distribution $\rho^V_{\pi,c}(x,a) = \mu^V_{\pi,c}(x)/A$. Let $\delta \in (0,1)$ and $\epsilon \in (0,1)$. Let $\overline{\beta}_u = \gamma \max_x \sum_{x'} P_{(x,\pi_*(x)),x'} u(x')/u(x)$ and

$$D = \frac{(1+\overline{\beta}_V)\mu^\top_{\pi_*,c}V}{2c^\top J_*} \sup_{w \in \mathcal{N}} \|J_* - \Psi w\|_{\infty,1/V} \ , \qquad m \geq \frac{16AD}{(1-\gamma)\epsilon}\left(d \log \frac{48AD}{(1-\gamma)\epsilon} + \log \frac{2}{\delta}\right) \ .$$

Let $\mathcal{S}$ be a set of m random state-action pairs sampled under $\rho^V_{\pi_*,c}$. Let $\widehat{w}$ be a solution of the following sampled LP:

$$\begin{aligned} &\max_{w \in \mathbb{R}^d} \ c^\top \Psi w \, , \\ &\text{s.t.} \quad w \in \mathcal{N}, \, \forall (x,a) \in \mathcal{S}, \, \ell(x,a) + \gamma P_{(x,a),:}\Psi w \geq (\Psi w)(x) \ . \end{aligned}$$

de Farias and Van Roy (2004) prove that with probability at least $1-\delta$, we have

$$\|J_* - \Psi \widehat{w}\|_{1,c} \leq \|J_* - \Psi w_*\|_{1,c} + \epsilon \|J_*\|_{1,c} \ .$$

This result has a number of limitations. First, vector $\mu_{\pi_*,c}$ (that is used in the definition of D) depends on the optimal policy, but an optimal policy is what we want to compute in the first place. Second, we can no longer use Inequality (8) to obtain a performance bound (a bound on $\|J_{\pi_{\Psi\widehat{w}}} - J_*\|_{1,c}$), as $\Psi\widehat{w}$ does not necessarily satisfy all constraints of ALP.

Desai et al. (2012) study a smoothed version of ALP, in which slack variables are introduced that allow for some violation of the constraints. Let D' be a violation budget. The smoothed ALP (SALP) has the form of

$$\begin{aligned} &\max_{w,s} c^\top \Psi w \, , \\ &\text{s.t.} \quad \Psi w \leq L\Psi w + s, \, \mu^\top_{\pi_*,c} s \leq D', \, s \geq \mathbf{0}, \end{aligned} \qquad\qquad \begin{aligned} &\max_{w,s} c^\top \Psi w - \frac{2\mu^\top_{\pi_*,c} s}{1-\gamma} \, , \\ &\text{s.t.} \quad \Psi w \leq L\Psi w + s, \, s \geq \mathbf{0} \ . \end{aligned}$$

The ALP on RHS is equivalent to LHS with a specific choice of D'. Let $\overline{U} = \{u \in \mathbb{R}^X \ : \ u \geq \mathbf{1}\}$ be a set of weight vectors. Desai et al. (2012) prove that if w_* is a solution to above problem, then

$$\|J_* - \Psi w_*\|_{1,c} \leq \inf_{w,u \in \overline{U}} \|J_* - \Psi w\|_{\infty,1/u} \left(c^\top u + \frac{2(\mu_{\pi_*,c}^\top u)(1+\beta_u)}{1-\gamma} \right) .$$

The above bound improves (9) as $\overline{U}$ is larger than U and RHS in the above bound is smaller than RHS of (9). Further, they prove that if η is a distribution and we choose $c = (1-\gamma)\eta^\top (I - \gamma P^{\pi_{\Psi w_*}})$, then

$$\left\| J_{\mu_{\Psi w_*}} - J_* \right\|_{1,\eta} \leq \frac{1}{1-\gamma} \left(\inf_{w,u \in \overline{U}} \|J_* - \Psi w\|_{\infty,1/u} \left(c^\top u + \frac{2(\mu_{\pi_*,\nu}^\top u)(1+\beta_u)}{1-\gamma} \right) \right) .$$

Similar methods are also proposed by Petrik and Zilberstein (2009). One problem with this result is that c is defined in terms of w_*, which itself depends on c. Also, the smoothed ALP formulation uses π_* which is not known. Desai et al. (2012) also propose a computationally efficient algorithm. Let $\mathcal{S}$ be a set of S random states drawn under distribution $\mu_{\pi_*,c}$. Let $\mathcal{N}' \subset \mathbb{R}^d$ be a known set that contains the solution of SALP. The algorithm solves the following LP:

$$\begin{aligned} &\max_{w,s} \ c^\top \Psi w - \frac{2}{(1-\gamma)S} \sum_{x \in \mathcal{S}} s(x) , \\ &\text{s.t.} \quad \forall x \in \mathcal{S}, \, (\Psi w)(x) \leq (L \Psi w)(x) + s(x), \, s \geq \mathbf{0}, \, w \in \mathcal{N}' . \end{aligned}$$

Let $\widehat{w}$ be the solution of this problem. Desai et al. (2012) prove high probability bounds on the approximation error $\|J_* - \Psi \widehat{w}\|_{1,c}$. However, it is no longer clear if a performance bound on $\|J_* - J_{\pi_{\Psi \widehat{w}}}\|_{1,c}$ can be obtained from this approximation bound.

Next, we turn our attention to average cost ALP, which is a more challenging problem and is also the focus of this paper. Let ν be a distribution over states, $u : \mathcal{X} \to [1, \infty)$, $\eta > 0$, $\gamma \in [0,1]$, $P_\gamma^\pi = \gamma P^\pi + (1-\gamma)\mathbf{1}\nu^\top$, and $L_\gamma h = \min_\pi (\ell_\pi + P_\gamma^\pi h)$. de Farias and Van Roy (2006) propose the following optimization problem:

$$\begin{aligned} &\min_{w,s_1,s_2} \ s_1 + \eta s_2 , \\ &\text{s.t.} \quad L_\gamma \Psi w - \Psi w + s_1 \mathbf{1} + s_2 u \geq \mathbf{0}, \, s_2 \geq 0 . \end{aligned} \tag{10}$$

Let $(w_*, s_{1,*}, s_{2,*})$ be the solution of this problem. Define the mixing time of policy π by

$$\tau_\pi = \inf \left\{ \tau \ : \ \left| \frac{1}{t} \sum_{t'=0}^{t-1} \nu^\top (P^\pi)^{t'} \ell_\pi - \lambda_\pi \right| \leq \frac{\tau}{t}, \forall t \right\} .$$

Let $\tau_* = \liminf_{\delta \to 0} \{\tau_\pi \ : \ \lambda_\pi \leq \lambda_* + \delta\}$. Let π_γ^* be the optimal policy when discount factor is

γ. Let $\pi_{\gamma,w}$ be the greedy policy with respect to Ψw when discount factor is γ, $\mu_{\gamma,\pi}^\top = (1-\gamma)\sum_{t=0}^{\infty}\gamma^t\nu^\top(P^\pi)^t$ and $\mu_{\gamma,w} = \mu_{\gamma,\pi_{\gamma,w}}$. de Farias and Van Roy (2006) prove that if $\eta \geq (2-\gamma)\mu_{\gamma,\pi_\gamma^*}^\top u$,

$$\lambda_{w_*} - \lambda_* \leq \frac{(1+\beta)\eta\max(D'',1)}{1-\gamma}\min_w \left\|h_\gamma^* - \Psi w\right\|_{\infty,1/u} + (1-\gamma)(\tau_* + \tau_{\pi_{w_*}}),$$

where $\beta = \max_\pi \|I - \gamma P^\pi\|_{\infty,1/u}$, $D'' = \mu_{\gamma,w_*}^\top V/(\nu^\top V)$ and $V = L_\gamma\Psi w_* - \Psi w_* + s_{1,*}\mathbf{1} + s_{2,*}u$. Similar results are obtained more recently by Veatch (2013).

An appropriate choice for vector ν is $\nu = \mu_{\gamma,w_*}$. Unfortunately, w_* depends on ν. We should also note that solving (10) can be computationally expensive. de Farias and Van Roy (2006) propose constraint sampling techniques similar to (de Farias and Van Roy, 2004), but no performance bounds are provided.

Wang et al. (2008) study ALP (3) and show that there is a dual form for standard value function based algorithms, including on-policy and off-policy updating and policy improvement. They also study the convergence of these methods, but no performance bounds are shown.

3 A Reduction to Stochastic Convex Optimization

In this section, we describe our algorithm as a reduction from Markov decision problems to stochastic convex optimization. The main idea is to convert the ALP (3) into an unconstrained optimization over Θ by adding a function of the constraint violations to the objective, then run stochastic gradient descent with unbiased estimated of the gradient.

For a positive constant H, form the following convex cost function by adding a multiple of the total constraint violations to the objective of the LP (3):

$$\begin{aligned} c(\theta) &= \ell^\top(\mu_0 + \Phi\theta) + H\left\|[\mu_0 + \Phi\theta]_-\right\|_1 + H\left\|(P-B)^\top(\mu_0+\Phi\theta)\right\|_1 \\ &= \ell^\top(\mu_0 + \Phi\theta) + H\left\|[\mu_0 + \Phi\theta]_-\right\|_1 + H\left\|(P-B)^\top\Phi\theta\right\|_1 \\ &= \ell^\top(\mu_0 + \Phi\theta) + H\sum_{(x,a)}\left|[\mu_0(x,a) + \Phi_{(x,a),:}\theta]_-\right| + H\sum_{x'}\left|(P-B)_{:,x'}^\top\Phi\theta\right| . \end{aligned} \tag{11}$$

We justify using this surrogate loss as follows. Suppose we find a near optimal vector $\widehat{\theta}$ such that $c(\widehat{\theta}) \leq \min_{\theta\in\Theta} c(\theta) + O(\epsilon)$. We will prove

1. that $\left\|[\mu_0 + \Phi\widehat{\theta}]_-\right\|_1$ and $\left\|(P-B)^\top(\mu_0+\Phi\widehat{\theta})\right\|_1$ are small and $\mu_0 + \Phi\widehat{\theta}$ is close to $\mu_{\widehat{\theta}}$ (by Lemma 2 in Section 3), and

2. that $\ell^\top(\mu_0 + \Phi\widehat{\theta}) \leq \min_{\theta\in\Theta} c(\theta) + O(\epsilon)$.

Input: Constant $S > 0$, number of rounds T, constant H.
Let Π_Θ be the Euclidean projection onto Θ.
Initialize $\theta_1 = 0$.
for $t := 1, 2, \ldots, T$ **do**
  Sample $(x_t, a_t) \sim q_1$ and $x'_t \sim q_2$.
  Compute subgradient estimate g_t (12).
  Update $\theta_{t+1} = \Pi_\Theta(\theta_t - \eta_t g_t)$.
end for
$\widehat{\theta}_T = \frac{1}{T}\sum_{t=1}^{T} \theta_t$.
Return policy $\pi_{\widehat{\theta}_T}$.

Figure 1: The Stochastic Subgradient Method for Markov Decision Processes

As we will show, these two facts imply that with high probability, for any $\theta \in \Theta$,

$$\mu_{\widehat{\theta}}^\top \ell \leq \mu_\theta^\top \ell + \frac{1}{\epsilon} \left\| [\mu_0 + \Phi\theta]_- \right\|_1 + \frac{1}{\epsilon} \left\| (P - B)^\top (\mu_0 + \Phi\theta) \right\|_1 + O(\epsilon) \ ,$$

which is to say that minimization of $c(\theta)$ solves the extended efficient large-scale ALP problem.

Unfortunately, calculating the gradients of $c(\theta)$ is $O(XA)$. Instead, we construct unbiased estimators and use stochastic gradient descent. Let T be the number of iterations of our algorithm. Let q_1 and q_2 be distributions over the state-action and state space, respectively (we will later discuss how to choose them). Let $((x_t, a_t))_{t=1\ldots T}$ be i.i.d. samples from q_1 and $(x'_t)_{t=1\ldots T}$ be i.i.d. samples from q_2. At round t, the algorithm estimates subgradient $\nabla c(\theta)$ by

$$g_t(\theta) = \ell^\top \Phi - H \frac{\Phi_{(x_t,a_t),:}}{q_1(x_t, a_t)} \mathbb{I}_{\{\mu_0(x_t,a_t) + \Phi_{(x_t,a_t),:}\theta < 0\}} + H \frac{(P - B)^\top_{:,x'_t} \Phi}{q_2(x'_t)} s((P - B)^\top_{:,x'_t} \Phi\theta). \tag{12}$$

This estimate is fed to the projected subgradient method, which in turn generates a vector θ_t. After T rounds, we average vectors $(\theta_t)_{t=1\ldots T}$ and obtain the final solution $\widehat{\theta}_T = \sum_{t=1}^{T} \theta_t / T$. Vector $\mu_0 + \Phi\widehat{\theta}_T$ defines a policy, which in turn defines a stationary distribution $\mu_{\widehat{\theta}_T}$.[2] The algorithm is shown in Figure 1.

[2]Recall that μ_θ is the stationary distribution of policy

$$\pi_\theta(a|x) = \frac{[\mu_0(x, a) + \Phi_{(x,a),:}\theta]_+}{\sum_{a'} [\mu_0(x, a') + \Phi_{(x,a'),:}\theta]_+} \ .$$

With an abuse of notation, we use μ_θ to denote policy π_θ as well.

3.1 Analysis

In this section, we state and prove our main result, Theorem 1. We begin with a discussion of the assumptions we make then follow with the main theorem. We break the proof into two main ingredients. First, we demonstrate that a good approximation to the surrogate loss gives a feature vector that is almost a stationary distribution; this is Lemma 2. Second, we justify the use of unbiased gradients in Theorem 3 and Lemma 5. The section concludes with the proof.

We make a mixing assumption on the MDP so that any policy quickly converges to its stationary distribution.

Assumption A1 ***(Fast Mixing)*** For any policy π, there exists a constant $\tau(\pi) > 0$ such that for all distributions d and d' over the state space, $\|dP^\pi - d'P^\pi\|_1 \leq e^{-1/\tau(\pi)} \|d - d'\|_1$.

Define

$$C_1 = \max_{(x,a)\in\mathcal{X}\times\mathcal{A}} \frac{\left\|\Phi_{(x,a),:}\right\|}{q_1(x,a)}, \qquad C_2 = \max_{x\in\mathcal{X}} \frac{\left\|(P-B)_{:,x}^\top \Phi\right\|}{q_2(x)} .$$

These constants appear in our performance bounds. So we would like to choose distributions q_1 and q_2 such that C_1 and C_2 are small. For example, if there is $C' > 0$ such that for any (x,a) and i, $\Phi_{(x,a),i} \leq C'/(XA)$ and each column of P has only N non-zero elements, then we can simply choose q_1 and q_2 to be uniform distributions. Then it is easy to see that

$$\frac{\left\|\Phi_{(x,a),:}\right\|}{q_1(x,a)} \leq C', \qquad \frac{\left\|(P-B)_{:,x}^\top \Phi\right\|}{q_2(x)} \leq C'(N+A) .$$

As another example, if $\Phi_{:,i}$ are exponential distributions and feature values at neighboring states are close to each other, then we can choose q_1 and q_2 to be appropriate exponential distributions so that $\left\|\Phi_{(x,a),:}\right\| / q_1(x,a)$ and $\left\|(P-B)_{:,x}^\top \Phi\right\| / q_2(x)$ are always bounded. Another example is when there exists a constant $C'' > 0$ such that for any x, $\left\|P_{:,x}^\top \Phi\right\| / \left\|B_{:,x}^\top \Phi\right\| < C''$[3] and we have access to an efficient algorithm that computes $Z_1 = \sum_{(x,a)} \left\|\Phi_{(x,a),:}\right\|$ and $Z_2 = \sum_x \left\|B_{:,x}^\top \Phi\right\|$ and can sample from $q_1(x,a) = \left\|\Phi_{(x,a),:}\right\| / Z_1$ and $q_2(x) = \left\|B_{:,x}^\top \Phi\right\| / Z_2$. In what follows, we assume that such distributions q_1 and q_2 are known.

We now state the main theorem.

Theorem 1. *Consider an expanded efficient large-scale dual ALP problem. Suppose we apply the stochastic subgradient method (shown in Figure 1) to the problem. Let $\epsilon \in (0,1)$. Let $T = 1/\epsilon^4$ be the number of rounds and $H = 1/\epsilon$ be the constraints multiplier in subgradient estimate* (12)*. Let $\widehat{\theta}_T$ be the output of the stochastic subgradient method after T rounds and let the learning rate be $\eta_1 = \cdots = \eta_T = S/(G'\sqrt{T})$, where $G' = \sqrt{d} + H(C_1 + C_2)$. Define $V_1(\theta) = \sum_{(x,a)} \left|[\mu_0(x,a) + \Phi_{(x,a),:}\theta]_-\right|$*

[3]This condition requires that columns of Φ are close to their *one step look-ahead*.

and $V_2(\theta) = \sum_{x'} \left|(P - B)^\top_{:,x'}(\mu_0 + \Phi\theta)\right|$. *Then, for any* $\delta \in (0,1)$, *with probability at least* $1-\delta$,

$$\mu_{\widehat{\theta}_T}^\top \ell \leq \min_{\theta \in \Theta} \left(\mu_\theta^\top \ell + \frac{1}{\epsilon}(V_1(\theta) + V_2(\theta)) + O(\epsilon) \right) , \tag{13}$$

where constants hidden in the big-O notation are polynomials in S, d, C_1, C_2, $\log(1/\delta)$, $V_1(\theta)$, $V_2(\theta)$, $\tau(\mu_\theta)$, *and* $\tau(\mu_{\widehat{\theta}_T})$.

Functions V_1 and V_2 are bounded by small constants for any set of normalized features: for any $\theta \in \Theta$,

$$\begin{aligned}
V_1(\theta) &\leq \|\mu_0\|_1 + \|\Phi\theta\|_1 \leq 1 + \sum_{(x,a)} \left|\Phi_{(x,a),:}\theta\right| \leq 1 + Sd\,, \\
V_2(\theta) &\leq \sum_{x'} \left|P^\top_{:,x'}(\mu_0 + \Phi\theta)\right| + \sum_{x'} \left|B^\top_{:,x'}(\mu_0 + \Phi\theta)\right| \\
&\leq \left(\sum_{x'} P_{:,x'}\right)^\top [\mu_0 + \Phi\theta]_+ + \left(\sum_{x'} B_{:,x'}\right)^\top [\mu_0 + \Phi\theta]_+ \\
&= (2)\mathbf{1}^\top [\mu_0 + \Phi\theta]_+ \\
&\leq (2)\mathbf{1}^\top |\mu_0 + \Phi\theta| \\
&= 2 + 2S\,.
\end{aligned}$$

Thus V_1 and V_2 can be very small given a carefully designed set of features. The output $\widehat{\theta}_T$ is a random vector as the algorithm is based on a stochastic convex optimization method. The above theorem shows that with high probability the policy implied by this output is near optimal.

The optimal choice for ϵ is $\epsilon = \sqrt{V_1(\theta_*) + V_2(\theta_*)}$, where θ_* is the minimizer of RHS of (13) and not known in advance. Once we obtain $\widehat{\theta}_T$, we can estimate $V_1(\widehat{\theta}_T)$ and $V_2(\widehat{\theta}_T)$ and use input $\epsilon = \sqrt{V_1(\widehat{\theta}_T) + V_2(\widehat{\theta}_T)}$ in a second run of the algorithm. This implies that the error bound scales like $O(\sqrt{V_1(\theta_*) + V_2(\theta_*)})$.

The next lemma, providing the first ingredient of the proof, relates the amount of constraint violation of a vector θ to resulting stationary distribution μ_θ.

Lemma 2. *Let* $u \in \mathbb{R}^{XA}$ *be a vector. Let* $\mathcal{N}$ *be the set of points* (x,a) *where* $u(x,a) < 0$ *and* $\mathcal{S}$ *be complement of* $\mathcal{N}$. *Assume*

$$\sum_{x,a} u(x,a) = 1, \sum_{(x,a)\in\mathcal{N}} |u(x,a)| \leq \epsilon', \left\|u^\top (P - B)\right\|_1 \leq \epsilon''.$$

Vector $[u]_+ / \|[u]_+\|_1$ *defines a policy, which in turn defines a stationary distribution* μ_u. *We have that*

$$\|\mu_u - u\|_1 \leq \tau(\mu_u)\log(1/\epsilon')(2\epsilon' + \epsilon'') + 3\epsilon'\,.$$

Proof. Let $f = u^\top(P - B)$. From $\left\|u^\top(P-B)\right\|_1 \le \epsilon''$, we get that for any $x' \in \mathcal{X}$,

$$\sum_{(x,a)\in\mathcal{S}} u(x,a)(P-B)_{(x,a),x'} = -\sum_{(x,a)\in\mathcal{N}} u(x,a)(P-B)_{(x,a),x'} + f(x')$$

such that $\sum_{x'} |f(x')| \le \epsilon''$. Let $h = [u]_+ / \|[u]_+\|_1$. Let $H' = \left\|h^\top(B-P)\right\|_1$. We write

$$\begin{aligned}
H' &= \sum_{x'} \left| \sum_{(x,a)\in\mathcal{S}} h(x,a)(B-P)_{(x,a),x'} \right| \\
&= \frac{1}{1+\epsilon'} \sum_{x'} \left| \sum_{(x,a)\in\mathcal{S}} u(x,a)(B-P)_{(x,a),x'} \right| \\
&= \frac{1}{1+\epsilon'} \sum_{x'} \left| -\sum_{(x,a)\in\mathcal{N}} u(x,a)(B-P)_{(x,a),x'} + f(x') \right| \\
&\le \frac{1}{1+\epsilon'} \left(\sum_{x'} \left| -\sum_{(x,a)\in\mathcal{N}} u(x,a)(B-P)_{(x,a),x'} \right| + \sum_{x'} |f(x')| \right) \\
&\le \frac{1}{1+\epsilon'} \left(\epsilon'' + \sum_{(x,a)\in\mathcal{N}} \sum_{x'} |u(x,a)| \left|(B-P)_{(x,a),x'}\right| \right) \\
&\le \frac{1}{1+\epsilon'} \left(\epsilon'' + \sum_{(x,a)\in\mathcal{N}} 2\,|u(x,a)| \right) \le \frac{2\epsilon' + \epsilon''}{1+\epsilon'} \\
&\le 2\epsilon' + \epsilon'' \,.
\end{aligned}$$

Vector h is almost a stationary distribution in the sense that

$$\left\|h^\top(B-P)\right\|_1 \le 2\epsilon' + \epsilon'' \,. \tag{14}$$

Let $\|w\|_{1,\mathcal{S}} = \sum_{(x,a)\in\mathcal{S}} |w(x,a)|$. First, we have that

$$\|h-u\|_1 \le \left\|h - \frac{u}{1+\epsilon'}\right\|_1 + \left\|u - \frac{u}{1+\epsilon'}\right\|_{1,\mathcal{S}} \le 2\epsilon' \,.$$

Next we bound $\|\mu_h - h\|_1$. Let $\nu_0 = h$ be the initial state distribution. We will show that as we run policy h (equivalently, policy μ_h), the state distribution converges to μ_h and this vector is close to h. From (14), we have $\mu_0^\top P = h^\top B + v_0$, where v_0 is such that $\|v_0\|_1 \le 2\epsilon' + \epsilon''$. Let M^h be a $X \times (XA)$ matrix that encodes policy h, $M^h_{(i,(i-1)A+1)\text{-}(i,iA)} = h(\cdot|x_i)$. Other entries of this matrix

are zero. We get that

$$h^\top P M^h = (h^\top B + v_0)M^h = h^\top B M^h + v_0 M^h = h^\top + v_0 M^h \,,$$

where we used the fact that $h^\top B M^h = h^\top$. Let $\mu_1^\top = h^\top P M^h$ which is the state-action distribution after running policy h for one step. Let $v_1 = v_0 M^h P = v_0 P^h$ and notice that as $\|v_0\|_1 \le 2\epsilon' + \epsilon''$, we also have that $\|v_1\|_1 = \left\|P^{h\top} v_0^\top\right\|_1 \le \|v_0\|_1 \le 2\epsilon' + \epsilon''$. Thus,

$$\mu_1^\top P = h^\top P + v_1 = h^\top B + v_0 + v_1 \,.$$

By repeating this argument for k rounds, we get that

$$\mu_k^\top = h^\top + (v_0 + v_1 + \cdots + v_{k-1})M^h$$

and it is easy to see that

$$\left\|(v_0 + v_1 + \cdots + v_{k-1})M^h\right\|_1 \le \sum_{i=0}^{k-1} \|v_i\|_1 \le k(2\epsilon' + \epsilon'').$$

Thus, $\|\mu_k - h\|_1 \le k(2\epsilon' + \epsilon'')$. Now notice that μ_k is the state-action distribution after k rounds of policy μ_h. Thus, by mixing assumption, $\|\mu_k - \mu_h\|_1 \le e^{-k/\tau(h)}$. By the choice of $k = \tau(h)\log(1/\epsilon')$, we get that $\|\mu_h - h\|_1 \le \tau(h)\log(1/\epsilon')(2\epsilon' + \epsilon'') + \epsilon'$.

□

The second ingredient is the validity of using estimates of the subgradients. We assume access to estimates of the subgradient of a convex cost function. Error bounds can be obtained from results in the stochastic convex optimization literature; the following theorem, a high-probability version of Lemma 3.1 of Flaxman et al. (2005) for stochastic convex optimization, is sufficient.

Theorem 3. *Let Z be a positive constant and $\mathcal{Z}$ be a bounded subset of $\mathbb{R}^d$ such that for any $z \in \mathcal{Z}$, $\|z\| \le Z$. Let $(f_t)_{t=1,2,\ldots,T}$ be a sequence of real-valued convex cost functions defined over $\mathcal{Z}$. Let $z_1, z_2, \ldots, z_T \in \mathcal{Z}$ be defined by $z_1 = 0$ and $z_{t+1} = \Pi_{\mathcal{Z}}(z_t - \eta f_t')$, where $\Pi_{\mathcal{Z}}$ is the Euclidean projection onto $\mathcal{Z}$, $\eta > 0$ is a learning rate, and $f_1', \ldots, f_T'$ are unbiased subgradient estimates such that $\mathbb{E}\left[f_t'|z_t\right] = \nabla f(z_t)$ and $\|f_t'\| \le F$ for some $F > 0$. Then, for $\eta = Z/(F\sqrt{T})$, for any $\delta \in (0,1)$, with probability at least $1 - \delta$,*

$$\sum_{t=1}^T f_t(z_t) - \min_{z\in\mathcal{Z}} \sum_{t=1}^T f_t(z) \le ZF\sqrt{T} + \sqrt{(1 + 4Z^2T)\left(2\log\frac{1}{\delta} + d\log\left(1 + \frac{Z^2T}{d}\right)\right)}\,. \tag{15}$$

Proof. Let $z_* = \operatorname{argmin}_{z\in\mathcal{Z}} \sum_{t=1}^T f_t(z)$ and $\eta_t = f_t' - \nabla f_t(z_t)$. Define function $h_t : \mathcal{Z} \to \mathbb{R}$ by $h_t(z) = f_t(z) + z\eta_t$. Notice that $\nabla h_t(z_t) = \nabla f_t(z_t) + \eta_t = f_t'$. By Theorem 1 of Zinkevich (2003),

we get that

$$\sum_{t=1}^{T} h_t(z_t) - \sum_{t=1}^{T} h_t(z_*) \le \sum_{t=1}^{T} h_t(z_t) - \min_{z\in\mathcal{Z}} \sum_{t=1}^{T} h_t(z) \le ZF\sqrt{T}\,.$$

Thus,

$$\sum_{t=1}^{T} f_t(z_t) - \sum_{t=1}^{T} f_t(z_*) \le ZF\sqrt{T} + \sum_{t=1}^{T} (z_* - z_t)\eta_t\,.$$

Let $S_t = \sum_{s=1}^{t-1}(z_* - z_s)\eta_s$, which is a self-normalized sum (de la Peña et al., 2009). By Corollary 3.8 and Lemma E.3 of Abbasi-Yadkori (2012), we get that for any $\delta \in (0,1)$, with probability at least $1-\delta$,

$$\begin{aligned}
|S_t| &\le \sqrt{\left(1 + \sum_{s=1}^{t-1}(z_t - z_*)^2\right)\left(2\log\frac{1}{\delta} + d\log\left(1 + \frac{Z^2 t}{d}\right)\right)} \\
&\le \sqrt{(1+4Z^2t)\left(2\log\frac{1}{\delta} + d\log\left(1 + \frac{Z^2 t}{d}\right)\right)}\,.
\end{aligned}$$

Thus,

$$\sum_{t=1}^{T} f_t(z_t) - \min_{z\in\mathcal{Z}} \sum_{t=1}^{T} f_t(z) \le ZF\sqrt{T} + \sqrt{(1+4Z^2T)\left(2\log\frac{1}{\delta} + d\log\left(1 + \frac{Z^2 T}{d}\right)\right)}\,.$$

□

Remark 4. *Let B_T denote RHS of* (15)*. If all cost functions are equal to f, then by convexity of f and an application of Jensen's inequality, we obtain that $f(\sum_{t=1}^{T} z_t/T) - \min_{z\in\mathcal{Z}} f(z) \le B_T/T$.*

As the next lemma shows, Theorem 3 can be applied in our problem to optimize cost function c.

Lemma 5. *Under the same conditions as in Theorem 1, we have that for any $\delta \in (0,1)$, with probability at least $1-\delta$,*

$$c(\widehat{\theta}_T) - \min_{\theta\in\Theta} c(\theta) \le \frac{SG'}{\sqrt{T}} + \sqrt{\frac{1+4S^2T}{T^2}\left(2\log\frac{1}{\delta} + d\log\left(1 + \frac{S^2T}{d}\right)\right)}\,. \tag{16}$$

Proof. We prove the lemma by showing that conditions of Theorem 3 are satisfied. We begin by calculating the subgradient and bounding its norm with a quantity independent of the number of states. If $\mu_0(x,a) + \Phi_{(x,a),:}\theta \ge 0$, then $\nabla_\theta \left|[\mu_0(x,a) + \Phi_{(x,a),:}\theta]_-\right| = 0$. Otherwise,

$\nabla_\theta \left|[\mu_0(x,a) + \Phi_{(x,a),:}\theta]_-\right| = -\Phi_{(x,a),:}$. Calculating,

$$\begin{aligned}\nabla_\theta c(\theta) &= \ell^\top \Phi + H \sum_{(x,a)} \nabla_\theta \left|[\mu_0(x,a) + \Phi_{(x,a),:}\theta]_-\right| + H \sum_{x'} \nabla_\theta \left|(P-B)^\top_{:,x'} \Phi\theta\right| \\ &= \ell^\top \Phi - H \sum_{(x,a)} \Phi_{(x,a),:} \mathbb{I}_{\{\mu_0(x,a)+\Phi_{(x,a),:}\theta<0\}} + H \sum_{x'} (P-B)^\top_{:,x'} \Phi s((P-B)^\top_{:,x'}\Phi\theta)\,,\end{aligned} \tag{17}$$

where $s(z) = \mathbb{I}_{\{z>0\}} - \mathbb{I}_{\{z<0\}}$ is the sign function. Let $\pm$ denote the plus or minus sign (the exact sign does not matter here). Let $G = \|\nabla_\theta c(\theta)\|$. We have that

$$G \le H \sqrt{\sum_{i=1}^d \left(\sum_{x'} \left(\pm \sum_{(x,a)} (P-B)_{(x,a),x'} \Phi_{(x,a),i}\right)\right)^2} + \left\|\ell^\top \Phi\right\| + H \sqrt{\sum_{i=1}^d \left(\sum_{(x,a)} \left|\Phi_{(x,a),i}\right|\right)^2}\,.$$

Thus,

$$\begin{aligned}G &\le \sqrt{\sum_{i=1}^d (\ell^\top \Phi_{:,i})^2} + H\sqrt{d} + H \sqrt{\sum_{i=1}^d \left(\sum_{(x,a)} \left(\pm \sum_{x'} (P-B)_{(x,a),x'}\right) \Phi_{(x,a),i}\right)^2} \\ &\le \sqrt{d} + H\sqrt{d} + H \sqrt{\sum_{i=1}^d \left(2 \sum_{(x,a)} \left|\Phi_{(x,a),i}\right|\right)^2} = \sqrt{d}(1+3H)\,,\end{aligned}$$

where we used $\left|\ell^\top \Phi_{:,i}\right| \le \|\ell\|_\infty \|\Phi_{:,i}\|_1 \le 1$.

Next, we show that norm of the subgradient estimate is bounded by G':

$$\|g_t\| \le \left\|\ell^\top \Phi\right\| + H \frac{\left\|\Phi_{(x_t,a_t),:}\right\|}{q_1(x_t,a_t)} + H \frac{\left\|(P-B)^\top_{:,x'_t} \Phi\right\|}{q_2(x'_t)} \le \sqrt{d} + H(C_1 + C_2)\,.$$

Finally, we show that the subgradient estimate is unbiased:

$$\begin{aligned}\mathbb{E}\left[g_t(\theta)\right] &= \ell^\top \Phi - H \sum_{(x,a)} q_1(x,a) \frac{\Phi_{(x,a),:}}{q_1(x,a)} \mathbb{I}_{\{\mu_0(x,a)+\Phi_{(x,a),:}\theta<0\}} \\ &\qquad + H \sum_{x'} q_2(x') \frac{(P-B)^\top_{:,x'}\Phi}{q_2(x')} s((P-B)^\top_{:,x'}\Phi\theta) \\ &= \ell^\top \Phi - H \sum_{(x,a)} \Phi_{(x,a),:} \mathbb{I}_{\{\mu_0(x,a)+\Phi_{(x,a),:}\theta<0\}} + H \sum_{x'} (P-B)^\top_{:,x'} \Phi s((P-B)^\top_{:,x'}\Phi\theta) \\ &= \nabla_\theta c(\theta)\,.\end{aligned}$$

The result then follows from Theorem 3 and Remark 4.

□

With both ingredients in place, we can prove our main result.

Proof of Theorem 1. Let b_T be the RHS of (16). Equation (16) implies that with high probability for any $\theta \in \Theta$,

$$\ell^\top(\mu_0 + \Phi\widehat{\theta}_T) + H\,V_1(\widehat{\theta}_T) + H\,V_2(\widehat{\theta}_T) \le \ell^\top(\mu_0 + \Phi\theta) + H\,V_1(\theta) + H\,V_2(\theta) + b_T\,. \tag{18}$$

From (18), we get that

$$V_1(\widehat{\theta}_T) \le \frac{1}{H}\left(2(1+S) + H\,V_1(\theta) + H\,V_2(\theta) + b_T\right) \stackrel{\text{def}}{=} \epsilon'\,, \tag{19}$$

$$V_2(\widehat{\theta}_T) \le \frac{1}{H}\left(2(1+S) + H\,V_1(\theta) + H\,V_2(\theta) + b_T\right) \stackrel{\text{def}}{=} \epsilon''\,. \tag{20}$$

Inequalities (19) and (20) and Lemma 2 give the following bound:

$$\left|\mu_{\widehat{\theta}_T}^\top \ell - (\mu_0 + \Phi\widehat{\theta}_T)^\top \ell\right| \le \tau(\mu_{\widehat{\theta}_T})\log(1/\epsilon')(2\epsilon' + \epsilon'') + 3\epsilon'\,. \tag{21}$$

From (18) we also have

$$\ell^\top(\mu_0 + \Phi\widehat{\theta}_T) \le \ell^\top(\mu_0 + \Phi\theta) + H\,V_1(\theta) + H\,V_2(\theta) + b_T\,,$$

which, together with (21) and Lemma 2, gives the final result:

$$\begin{aligned}
\mu_{\widehat{\theta}_T}^\top \ell &\le \ell^\top(\mu_0 + \Phi\theta) + H\,V_1(\theta) + H\,V_2(\theta) + b_T + \tau(\mu_{\widehat{\theta}_T})\log(1/\epsilon')(2\epsilon' + \epsilon'') + 3\epsilon' \\
&\le \mu_\theta^\top \ell + H\,V_1(\theta) + H\,V_2(\theta) + b_T + \tau(\mu_{\widehat{\theta}_T})\log(1/\epsilon')(2\epsilon' + \epsilon'') + 3\epsilon' \\
&\quad + \tau(\mu_\theta)\log(1/V_1(\theta))(2V_1(\theta) + V_2(\theta)) + 3V_1(\theta)\,.
\end{aligned}$$

Recall that $b_T = O(H/\sqrt{T})$. Because $H = 1/\epsilon$ and $T = 1/\epsilon^4$, we get that with high probability, for any $\theta \in \Theta$, $\mu_{\widehat{\theta}_T}^\top \ell \le \mu_\theta^\top \ell + \frac{1}{\epsilon}(V_1(\theta) + V_2(\theta)) + O(\epsilon)$.

□

Let's compare Theorem 1 with results of de Farias and Van Roy (2006). Their approach is to relate the original MDP to a perturbed version[4] and then analyze the corresponding ALP. (See Section 2 for more details.) Let Ψ be a feature matrix that is used to estimate value functions. Recall that λ_* is the average loss of the optimal policy and λ_w is the average loss of the greedy policy with respect to value function Ψw. Let h_γ^* be the differential value function when the restart probability in the perturbed MDP is $1-\gamma$. For vector v and positive vector u, define the

[4]In a perturbed MDP, the state process restarts with a certain probability to a *restart distribution*. Such perturbed MDPs are closely related to discounted MDPs.

weighted maximum norm $\|v\|_{\infty,u} = \max_x u(x)\,|v(x)|$. de Farias and Van Roy (2006) prove that for appropriate constants $C, C' > 0$ and weight vector u,

$$\lambda_{w_*} - \lambda_* \leq \frac{C}{1-\gamma} \min_w \left\| h^*_\gamma - \Psi w \right\|_{\infty,u} + C'(1-\gamma)\,. \tag{22}$$

This bound has similarities to bound (13): tightness of both bounds depends on the quality of feature vectors in representing the relevant quantities (stationary distributions in (13) and value functions in (22)). Once again, we emphasize that the algorithm proposed by de Farias and Van Roy (2006) is computationally expensive and requires access to a distribution that depends on optimal policy.

4 Sampling Constraints

In this section we describe our second algorithm for average cost MDP problems. Using the results on polytope constraint sampling (de Farias and Van Roy, 2004, Calafiore and Campi, 2005, Campi and Garatti, 2008), we reduce approximate the solution to the dual ALP with the solution to a smaller, sampled LP. Basically, de Farias and Van Roy (2004) claim that given a set of affine constraints in $\mathbb{R}^d$ and some measure q over these constraints, if we sample $k = O(d\log(1/\delta)/\epsilon)$ constraints, then with probability at least $1-\delta$, any point that satisfies all of these k sampled constraints also satisfies $1-\epsilon$ of the original set of constraints under measure q. This result is independent of the number of original constraints.

Let $\mathcal{L}$ be a family of affine constraints indexed by i: constraint i is satisfied at point $w \in \mathbb{R}^d$ if $a_i^\top w + b_i \geq 0$. Let $\mathcal{I}$ be the family of constraints by selecting k random constraints in $\mathcal{L}$ with respect to measure q.

Theorem 6 (de Farias and Van Roy (2004))**.** *Assume there exists a vector that satisfies all constraints in $\mathcal{L}$. For any δ and ϵ, if we take $m \geq \frac{4}{\epsilon}\left(d\log\frac{12}{\epsilon} + \log\frac{2}{\delta}\right)$, then, with probability $1-\delta$, a set $\mathcal{I}$ of m i.i.d. random variables drawn from $\mathcal{L}$ with respect to distribution q satisfies*

$$\sup_{\{w:\forall i\in\mathcal{I}, a_i^\top w + b_i \geq 0\}} q(\{j : a_j^\top w + b_j < 0\}) \leq \epsilon\,.$$

Our algorithm takes the following inputs: a positive constant S, a stationary distribution μ_0, a set $\Theta = \{\theta \;:\; \theta^\top \Phi^\top \mathbf{1} = 1 - \mu_0^\top \mathbf{1},\, \|\theta\| \leq S\}$, a distribution q_1 over the state-action space, a distribution q_2 over the state space, and constraint violation functions $v_1 : \mathcal{X}\times\mathcal{A} \to [-1, 0]$ and $v_2 : \mathcal{X} \to [0, 1]$. We will consider two families of constraints:

$$\mathcal{L}_1 = \{\mu_0(x,a) + \Phi_{(x,a),:}\theta \geq v_1(x,a) \mid (x,a) \in \mathcal{X}\times\mathcal{A}\}\,,$$
$$\mathcal{L}_2 = \left\{(P-B)^\top_{:,x}(\mu_0 + \Phi\theta) \leq v_2(x) \mid x \in \mathcal{X}\right\} \bigcup \left\{(P-B)^\top_{:,x}(\mu_0 + \Phi\theta) \geq -v_2(x) \mid x \in \mathcal{X}\right\}\,.$$

Input: Constant $S > 0$, stationary distribution μ_0, distributions q_1 and q_2, constraint violation functions v_1 and v_2, number of samples k_1 and k_2.
For $i = 1, 2$, let $\mathcal{I}_i$ be k_i constraints sampled from $\mathcal{L}_i$ under distribution q_i.
Let $\mathcal{I}$ be the set of vectors that satisfy all constraints in $\mathcal{I}_1$ and $\mathcal{I}_2$.
Let $\widetilde{\theta}$ be the solution to LP:

$$\begin{aligned} &\min_{\theta \in \Theta} \ell^\top (\mu_0 + \Phi\theta)\,, \\ &\text{s.t.} \quad \theta \in \mathcal{I},\, \theta \in \Theta\,. \end{aligned} \tag{24}$$

Return policy $\mu_{\widetilde{\theta}}$.

Figure 2: The Constraint Sampling Method for Markov Decision Processes

Let θ_* be the solution of

$$\begin{aligned} &\min_{\theta \in \Theta} \ell^\top (\mu_0 + \Phi\theta)\,, \\ &\text{s.t.} \quad \theta \in \mathcal{L}_1,\, \theta \in \mathcal{L}_2,\, \theta \in \Theta\,. \end{aligned} \tag{23}$$

The constraint sampling algorithm is shown in Figure 2. We refer to (24) as the sampled ALP, while we refer to (3) as the full ALP.

4.1 Analysis

We require Assumption A1 as well as:

Assumption A2 ***(Feasibility)*** There exists a vector that satisfies all constraints $\mathcal{L}_1$ and $\mathcal{L}_2$.

Validity of this assumption depends on the choice of functions v_1 and v_2. Larger functions ensure that this assumption is satisfied, but as we show, this leads to larger error.

The next two lemmas apply theorem 6 to constraints $\mathcal{L}_1$ and $\mathcal{L}_2$, respectively.

Lemma 7. *Let $\delta_1 \in (0,1)$ and $\epsilon_1 \in (0,1)$. If we choose $k_1 = \frac{4}{\epsilon_1}\left(d \log \frac{12}{\epsilon_1} + \log \frac{2}{\delta_1}\right)$, then with probability at least $1 - \delta_1$, $\sum_{(x,a)} \left|[\mu_0(x,a) + \Phi_{(x,a),:}\widetilde{\theta}]_-\right| \leq SC_1\epsilon_1 + \|v_1\|_1$.*

Proof. Applying theorem 6, we have that w.p. $1 - \delta_1$, $q_1(\mu_0(x,a) + \Phi_{(x,a),:}\widetilde{\theta} \geq v_1(x,a)) \geq 1 - \epsilon_1$, and thus

$$\sum_{(x,a)} q_1(x,a) \mathbb{I}_{\{\mu_0(x,a) + \Phi_{(x,a),:}\widetilde{\theta} < v_1(x,a)\}} \leq \epsilon_1\,.$$

Let $L = \sum_{(x,a)} \left| [\mu_0(x,a) + \Phi_{(x,a),:}\widetilde{\theta}]_- \right|$. With probability $1 - \delta_1$,

$$\begin{aligned}
L &= \sum_{(x,a)} \left| [\mu_0(x,a) + \Phi_{(x,a),:}\widetilde{\theta}]_- \right| \mathbb{I}_{\{\mu_0(x,a)+\Phi_{(x,a),:}\widetilde{\theta} \le v_1(x,a)\}} \\
&\quad + \sum_{(x,a)} \left| [\mu_0(x,a) + \Phi_{(x,a),:}\widetilde{\theta}]_- \right| \mathbb{I}_{\{\mu_0(x,a)+\Phi_{(x,a),:}\widetilde{\theta} > v_1(x,a)\}} \\
&\le \sum_{(x,a)} \left| \Phi_{(x,a),:}\widetilde{\theta} \right| \mathbb{I}_{\{\mu_0(x,a)+\Phi_{(x,a),:}\widetilde{\theta} \le v_1(x,a)\}} + \|v_1\|_1 \\
&\le \sum_{(x,a)} \left\| \Phi_{(x,a),:} \right\| \left\| \widetilde{\theta} \right\| \mathbb{I}_{\{\mu_0(x,a)+\Phi_{(x,a),:}\widetilde{\theta} \le v_1(x,a)\}} + \|v_1\|_1 \\
&\le \sum_{(x,a)} SC_1 q_1(x,a) \mathbb{I}_{\{\mu_0(x,a)+\Phi_{(x,a),:}\widetilde{\theta} \le v_1(x,a)\}} + \|v_1\|_1 \\
&\le SC_1\epsilon_1 + \|v_1\|_1 \ .
\end{aligned}$$

□

Lemma 8. *Let $\delta_2 \in (0,1)$ and $\epsilon_2 \in (0,1)$. If we choose $k_2 = \frac{4}{\epsilon_2}\left(d \log \frac{12}{\epsilon_2} + \log \frac{2}{\delta_2}\right)$, then with probability at least $1 - \delta_2$, $\left\| (P-B)^\top \Phi\widetilde{\theta} \right\|_1 \le SC_2\epsilon_2 + \|v_2\|_1$.*

Proof. Applying theorem 6, we have that $q_2\left(\left|(P-B)_{:,x}^\top \Phi\widetilde{\theta}\right| \le v_2(x)\right) \ge 1 - \epsilon_2$. This yields

$$\sum_x q_2(x) \mathbb{I}_{\{\left|(P-B)_{:,x}^\top \Phi\widetilde{\theta}\right| \ge v_2(x)\}} \le \epsilon_2 \ . \tag{25}$$

Let $L' = \sum_x \left|(P-B)_{:,x}^\top \Phi\widetilde{\theta}\right|$. Thus, with probability $1 - \delta_2$,

$$\begin{aligned}
L' &= \sum_x \left|(P-B)_{:,x}^\top \Phi\widetilde{\theta}\right| \mathbb{I}_{\{\left|(P-B)_{:,x}^\top \Phi\widetilde{\theta}\right| > v_2(x)\}} \\
&\quad + \sum_x \left|(P-B)_{:,x}^\top \Phi\widetilde{\theta}\right| \mathbb{I}_{\{\left|(P-B)_{:,x}^\top \Phi\widetilde{\theta}\right| \le v_2(x)\}} \\
&\le \sum_x \left\|(P-B)_{:,x}^\top \Phi\right\| \left\|\widetilde{\theta}\right\| \mathbb{I}_{\{\left|(P-B)_{:,x}^\top \Phi\widetilde{\theta}\right| > v_2(x)\}} + \|v_2\|_1 \\
&\le \sum_x SC_2 q_2(x) \mathbb{I}_{\{\left|(P-B)_{:,x}^\top \Phi\widetilde{\theta}\right| > v_2(x)\}} + \|v_2\|_1 \\
&\le SC_2\epsilon_2 + \|v_2\|_1 \ ,
\end{aligned}$$

where the last step follows from (25).

□

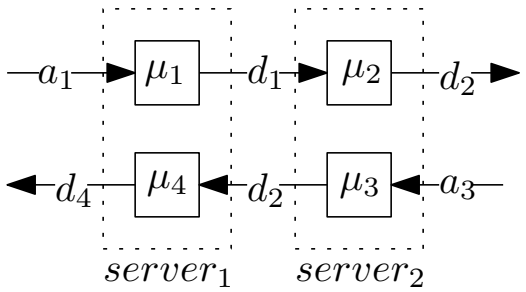

Figure 3: The 4D queueing network. Customers arrive at queue μ_1 or μ_3 then are referred to queue μ_2 or μ_4, respectively. Server 1 can either process queue 1 or 4, and server 2 can only process queue 2 or 3.

We are ready to prove the main result of this section. Let $\widetilde{\theta}$ denote the solution of the sampled ALP, θ_* denote the solution of the full ALP (23), and $\mu_{\widetilde{\theta}}$ be the stationary distribution of the solution policy. Our goal is to compare $\ell^\top \mu_{\widetilde{\theta}}$ and $\ell^\top \mu_{\theta_*}$.

Theorem 9. *Let $\epsilon \in (0,1)$ and $\delta \in (0,1)$. Let $\epsilon' = SC_1\epsilon + \|v_1\|_1$ and $\epsilon'' = SC_2\epsilon + \|v_2\|_1$. If we sample constraints with $k_1 = \frac{4}{\epsilon}\left(d\log\frac{12}{\epsilon} + \log\frac{4}{\delta}\right)$ and $k_2 = \frac{4}{\epsilon}\left(d\log\frac{12}{\epsilon} + \log\frac{4}{\delta}\right)$, then, with probability $1-\delta$,*

$$\begin{aligned}\ell^\top \mu_{\widetilde{\theta}} \le \ell^\top \mu_{\theta_*} &+ \tau(\mu_{\widetilde{\theta}})\log(1/\epsilon')(2\epsilon' + \epsilon'') + 3\epsilon' \\ &+ \tau(\mu_*)\log(1/\|v_1\|)(2\|v_1\| + \|v_2\|) + 3\|v_1\| \ .\end{aligned}$$

Proof. Let $\delta_1 = \delta_2 = \delta/2$. By Lemmas 7 and 8, w.p. $1-\delta$, $\sum_{(x,a)} \left|[\mu_0(x,a) + \Phi_{(x,a),:}\widetilde{\theta}]_-\right| \le \epsilon'$ and $\left\|(P-B)^\top(\mu_0 + \Phi\widetilde{\theta})\right\|_1 \le \epsilon''$. Then by Lemma 2,

$$\left|\ell^\top \mu_{\widetilde{\theta}} - \ell^\top(\mu_0 + \Phi\widetilde{\theta})\right| \le \tau(\mu_{\widetilde{\theta}})\log(1/\epsilon')(2\epsilon' + \epsilon'') + 3\epsilon' \ .$$

We also have that $\ell^\top(\mu_0 + \Phi\widetilde{\theta}) \le \ell^\top(\mu_0 + \Phi\theta_*)$. Thus,

$$\begin{aligned}\ell^\top \mu_{\widetilde{\theta}} &\le \ell^\top(\mu_0 + \Phi\theta_*) + \tau(\mu_{\widetilde{\theta}})\log(1/\epsilon')(2\epsilon' + \epsilon'') + 3\epsilon' \\ &\le \ell^\top \mu_{\theta_*} + \tau(\mu_{\widetilde{\theta}})\log(1/\epsilon')(2\epsilon' + \epsilon'') + 3\epsilon' \\ &\quad + \tau(\mu_{\theta_*})\log(1/\|v_1\|)(2\|v_1\| + \|v_2\|) + 3\|v_1\|\,,\end{aligned}$$

where the last step follows from Lemma 2. ∎

5 Experiments

In this section, we apply both algorithms to the four-dimensional discrete-time queueing network illustrated in Figure 5. This network has a relatively long history; see, e.g. Rybko and Stolyar (1992) and more recently de Farias and Van Roy (2003a) (c.f. section 6.2). There are four queues, $\mu_1, \ldots, \mu_4$, each with state $0, \ldots, B$. Since the cardinality of the state space is $X = (1+B)^4$, even a modest B results in huge state-spaces. For time t, let $X_t \in X$ be the state and $s_{i,t} \in \{0, 1\}$, $i = 1, 2, 3, 3$ denote whether queue i is being served. Server 1 only serves queue 1 or 4, server 2 only serves queue 2 or 3, and neither server can idle. Thus, $s_{1,t} + s_{4,t} = 1$ and $s_{2,t} + s_{3,t} = 1$. The dynamics are as follows. At each time t, the following random variables are sampled independently: $A_{1,t} \sim \text{Bernoulli}(a_1)$, $A_{3,t} \sim \text{Bernoulli}(a_3)$, and $D_{i,t} \sim \text{Bernoulli}(d_i * s_{i,t})$ for $i = 1, 2, 3, 4$. Using $e_1, \ldots, e_4$ to denote the standard basis vectors, the dynamics are:

$$\begin{aligned} X'_{t+1} =& X_t + A_{1,t}e_1 + A_{3,t}e_3 \\ &+ D_{1,t}(e_2 - e_1) - D_{2,t}e_2 \\ &+ D_{3,t}(e_4 - e_3) - D_{4,t}e_4, \end{aligned}$$

and $X_{t+1} = \max(\mathbf{0}, \min((B), X'_{t+1}))$ (i.e. all four states are thresholded from below by 0 and above by B). The loss function is the total queue size: $\ell(X_t) = ||X_t||_1$. We compared our method against two common heuristics. In the first, denoted LONGER, each server operates on the queue that is longer with ties broken uniformly at random (e.g. if queue 1 and 4 had the same size, they are equally likely to be served). In the second, denoted LBFS (last buffer first served), the downstream queues always have priority (server 1 will serve queue 4 unless it has length 0, and server 2 will serve queue 2 unless it has length 0). These heuristics are common and have been used an benchmarks for queueing networks (e.g. de Farias and Van Roy (2003a)).

We used $a_1 = a_3 = .08$, $d_1 = d_2 = .12$, and $d_3 = d_4 = .28$, and buffer sizes $B_1 = B_4 = 38$, $B_2 = B_3 = 25$ as the parameters of the network.. The asymmetric size was chosen because server 1 is the bottleneck and tend to have has longer queues. The first two features are the stationary distributions corresponding to two heuristics. We also included two types of non-stationary-distribution features. For every interval $(0, 5], (6, 10], \ldots, (45, 50]$ and action A, we added a feature ψ with $\varphi(x, a) = 1$ if $\ell(x, a)$ is in the interval and $a = A$. To define the second type, consider the three intervals $I_1 = [0, 10]$, $I_2 = [11, 20]$, and $I_3 = [21, 25]$. For every 4-tuple of intervals $(J_1, J_2, J_3, J_4) \in \{I_1, I_2, I_3\}^4$ and action A, we created a feature ψ with $\psi(x, a) = 1$ only if $x_i \in J_i$ and $a = A$. Every feature was normalized to sum to 1. In total, we had 372 features which is about a 10^4 reduction in dimension from the original problem.

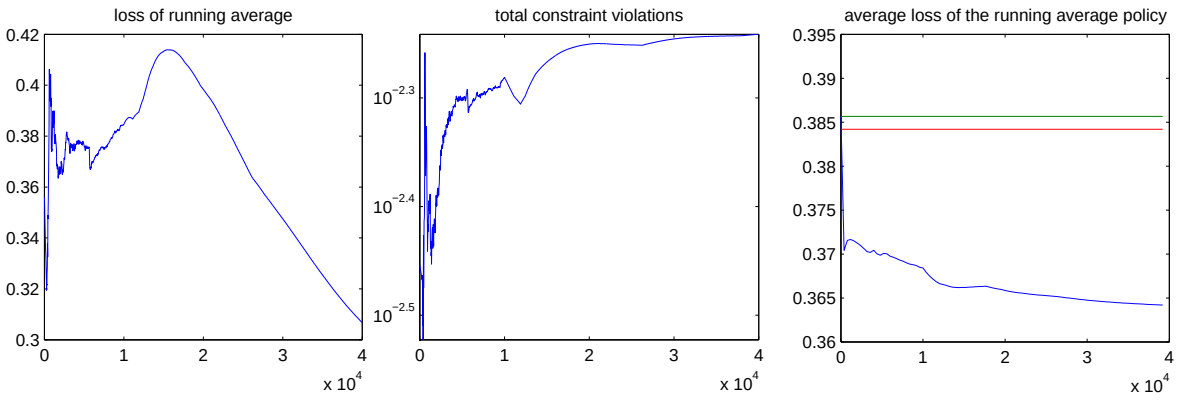

Figure 4: The left plot is of the linear objective of the running average, i.e. $\ell^\top \Phi \widehat{\theta}_t$. The center plot is the sum of the two constraint violations of $\widehat{\theta}_t$, and the right plot is $\ell^\top \tilde{\mu}_{\widehat{\theta}_t}$ (the average loss of the derived policy). The two horizontal lines correspond to the loss of the two heuristics, LONGER and LBFS.

5.1 Stochastic Gradient Descent

We ran our stochastic gradient descent algorithm with $I = 1000$ sampled constraints and constraint gain $H = 2$. Our learning rate began at 10^{-4} and halved every 2000 iterations. The results of our algorithm are plotted in Figure 5.1, where $\widehat{\theta}_t$ denotes the running average of θ_t. The left plot is of the LP objective, $\ell^\top(\mu_0 + \Phi\widehat{\theta}_t)$. The middle plot is of the sum of the constraint violations, $\left\|[\mu_0 + \Phi\widehat{\theta}_t]_-\right\|_1 + \left\|(P-B)^\top \Phi\widehat{\theta}_t\right\|_1$. Thus, $c(\widehat{\theta}_t)$ is a scaled sum of the first two plots. Finally, the right plot is of the average losses, $\ell^\top \mu_{\widehat{\theta}_t}$ and the two horizontal lines correspond to the loss of the two heuristics, LONGER and LBFS. The right plot demonstrates that, as predicted by our theory, minimizing the surrogate loss $c(\theta)$ does lead to lower average losses.

All previous algorithms (including de Farias and Van Roy (2003a)) work with value functions, while our algorithm works with stationary distributions. Due to this difference, we cannot use the same feature vectors to make a direct comparison. The solution that we find in this different approximating set is slightly worse than the solution of de Farias and Van Roy (2003a).

5.2 Constraint Sampling

For the constraint sampling algorithm, we sampled the simplex constraints uniformly with 10 different sample sizes: 508, 792, 1235, 1926, 3003, 4684, 7305, 11393, 17768, and 27712. Since $XA = 4.1 * 10^6$, these sample sizes correspond to less that 1%. The stationary constraints were sampled in the same proportion (i.e. A times fewer samples). Let $a_1, \ldots, a_{AN}$ and $b_1, \ldots, b_N$ be the indices of the sampled simplex and stationary constraints, respectively. Explicitly, the sampled

LP is:

$$\min_{\theta} (\Phi\theta)^\top \ell\,, \tag{26}$$

$$\text{s.t.} \quad (\Phi\theta)^\top \mathbf{1} = 1,\ \Phi_{a_i,:}\theta \geq \mathbf{0},\ \forall i = 1, \ldots, AN,$$

$$\left|\Phi\theta^\top (P - B)_{:,b_i}\right| \leq \epsilon_s,\ \forall i = 1, \ldots, N,\ \|\theta\|_\infty \leq M \tag{27}$$

where M and ϵ are necessary to ensure the LP always has a feasible and bounded solution. This corresponds to setting $v_1 = 0$ and $v_2 = \epsilon$. In particular, we used $M = 3$ and $\epsilon = 10^{-3}$. Using differenc values of ϵ did not have a large effect on the behavior of the algorithm.

For each sample size, we sample the constraints, solve the LP, then simulate the average loss of the policy. We repeated this procedure 35 times for each sample size and plotted the mean with error bars corresponding to the variance across each sample size in Figure 5.2. Note the log scale on the x-axis. The best loss corresponds to 4684 sampled simplex constraints, or roughly 1%, and is a marked improvement over the average loss found by the stochastic gradient descent method. However, changing the sample size by a factor of 4 in either direction is enough to obliterate this advantage.

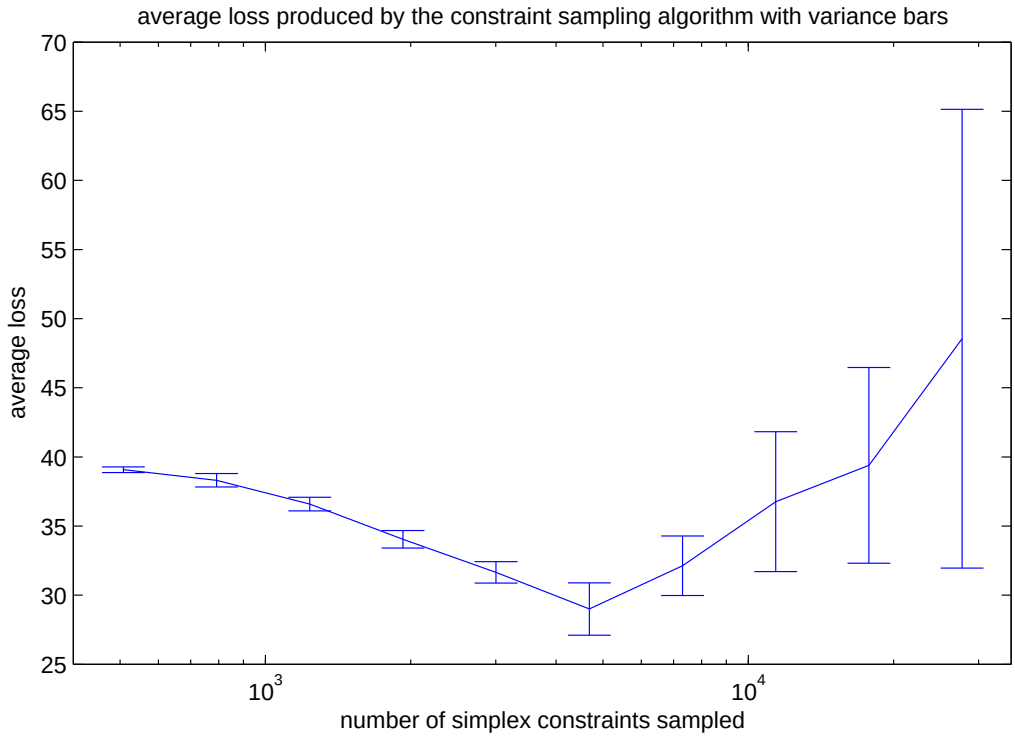

Figure 5: Average loss with variance error bars of the constraint sampling algorithm run with a variety of sample sizes.

First, we notice that the mean average loss is not monotonic. If we use too few constraints, then the sampled ALP does not reflect our original problem and we expect that the solution will make poor policies. On the other hand, if we sample too many constraints, then the LP is too restrictive and cannot adequately explore the feature space. To explain the increasing variance, recall that we have three families of constraints: the simplex constraints, the stationarity constraints, and the box constraints (i.e. $|\theta|_\infty \leq M$). Only the simplex and stationarity constraints are sampled. For the small sample sizes, the majority of the active constraints are the box constraint so $\tilde{\theta}$ (the minimizer of the LP) is not very sensitive to the random sample. However, as the sample size grows, so does the number of active simplex and stationarity constraints; hence, the random constraints affect $\tilde{\theta}$ to a greater degree and the variance increases.

6 Conclusions

In this paper, we defined and solved the extended large-scale efficient ALP problem. We proved that, under certain assumptions about the dynamics, the stochastic subgradient method produces a policy with average loss competitive to all $\theta \in \Theta$, not just all θ producing a stationary distribution. We demonstrated this algorithm on the Rybko-Stoylar four-dimensional queueing network and recovered a policy better than two common heuristics and comparable to previous results on ALPs de Farias and Van Roy (2003a). A future direction is to find other interesting regularity conditions under which we can handle large-scale MDP problems. We also plan to conduct more experiments with challenging large-scale problems.

7 Acknowledgements

We gratefully acknowledge the support of the NSF through grant CCF-1115788 and of the ARC through an Australian Research Council Australian Laureate Fellowship (FL110100281).

References

Y. Abbasi-Yadkori. *Online Learning for Linearly Parametrized Control Problems*. PhD thesis, University of Alberta, 2012.

R. Bellman. *Dynamic Programming*. Princeton University Press, 1957.

D. P. Bertsekas. *Dynamic Programming and Optimal Control*. Athena Scientific, 2007.

D. P. Bertsekas and J. Tsitsiklis. *Neuro-Dynamic Programming*. Athena scientific optimization and computation series. Athena Scientific, 1996.

G. Calafiore and M. C. Campi. Uncertain convex programs: randomized solutions and confidence levels. *Mathematical Programming*, 102(1):25–46, 2005.

M. C. Campi and S. Garatti. The exact feasibility of randomized solutions of uncertain convex programs. *SIAM Journal on Optimization*, 19(3):1211–1230, 2008.

D. P. de Farias and B. Van Roy. The linear programming approach to approximate dynamic programming. *Operations Research*, 51, 2003a.

D. P. de Farias and B. Van Roy. Approximate linear programming for average-cost dynamic programming. In *Advances in Neural Information Processing Systems (NIPS)*, 2003b.

D. P. de Farias and B. Van Roy. On constraint sampling in the linear programming approach to approximate dynamic programming. *Mathematics of Operations Research*, 29, 2004.

D. P. de Farias and B. Van Roy. A cost-shaping linear program for average-cost approximate dynamic programming with performance guarantees. *Mathematics of Operations Research*, 31, 2006.

V. H. de la Peña, T. L. Lai, and Q-M. Shao. *Self-normalized processes: Limit theory and Statistical Applications*. Springer, 2009.

V. V. Desai, V. F. Farias, and C. C. Moallemi. Approximate dynamic programming via a smoothed linear program. *Operations Research*, 60(3):655–674, 2012.

A. D. Flaxman, A. T. Kalai, and H. B. McMahan. Online convex optimization in the bandit setting: gradient descent without a gradient. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, 2005.

C. Guestrin, M. Hauskrecht, and B. Kveton. Solving factored mdps with continuous and discrete variables. In *Twentieth Conf. Uncertainty in Artificial Intelligence*, 2004.

M. Hauskrecht and B. Kveton. Linear program approximations to factored continuous-state markov decision processes. In *Advances in Neural Information Processing Systems*, 2003.

R. A. Howard. *Dynamic Programming and Markov Processes*. MIT, 1960.

H. R. Maei, Cs. Szepesvári, S. Bhatnagar, D. Precup, D. Silver, and R. S. Sutton. Convergent temporal-difference learning with arbitrary smooth function approximation. In *Advances in Neural Information Processing Systems*, 2009.

H. R. Maei, Cs. Szepesvári, S. Bhatnagar, and R. S. Sutton. Toward off-policy learning control with function approximation. In *Proceedings of the 27th International Conference on Machine Learning*, 2010.

A. S. Manne. Linear programming and sequential decisions. *Management Science*, 6(3):259–267, 1960.

M. Petrik and S. Zilberstein. Constraint relaxation in approximate linear programs. In *Proc. 26th Internat. Conf. Machine Learning (ICML)*, 2009.

A. N. Rybko and A. L. Stolyar. Ergodicity of stochastic processes describing the operation of open queueing networks. *Problemy Peredachi Informatsii*, 28(3):3–26, 1992.

P. Schweitzer and A. Seidmann. Generalized polynomial approximations in Markovian decision processes. *Journal of Mathematical Analysis and Applications*, 110:568–582, 1985.

R. S. Sutton and A. G. Barto. *Reinforcement Learning: An Introduction*. Bradford Book. MIT Press, 1998.

R. S. Sutton, H. R. Maei, D. Precup, S. Bhatnagar, D. Silver, Cs. Szepesvári, and E. Wiewiora. Fast gradient-descent methods for temporal-difference learning with linear function approximation. In *Proceedings of the 26th International Conference on Machine Learning*, 2009a.

R. S. Sutton, Cs. Szepesvári, and H. R. Maei. A convergent O(n) algorithm for off-policy temporal-difference learning with linear function approximation. In *Advances in Neural Information Processing Systems*, 2009b.

M. H. Veatch. Approximate linear programming for average cost mdps. *Mathematics of Operations Research*, 38(3), 2013.

T. Wang, D. Lizotte, M. Bowling, and D. Schuurmans. Dual representations for dynamic programming. *Journal of Machine Learning Research*, pages 1–29, 2008.

M. Zinkevich. Online convex programming and generalized infinitesimal gradient ascent. In *ICML*, 2003.